arXiv:1702.05000v1 [cs.CR] 16 Feb 2017

Existential Assertions for Voting Protocols

R. Ramanujam
Institute of Mathematical Sciences
Chennai, India.
jam@imsc.res.in

Vaishnavi Sundararajan
Chennai Mathematical Institute, Chennai, India.
vaishnavi@cmi.ac.in

S.P. Suresh
Chennai Mathematical Institute, Chennai, India.
spsuresh@cmi.ac.in

Abstract

In [22], we extend the Dolev-Yao model with assertions. We build on that work and add existential abstraction to the language, which allows us to translate common constructs used in voting protocols into proof properties. We also give an equivalence-based definition of anonymity in this model, and prove anonymity for the FOO voting protocol.

1 Anonymity

Formal verification of security protocols often involves the analysis of a property where the relationship between an agent and a message sent by him/her needs to be kept secret. This property, called "anonymity", is a version of the general unlinkability property, and one of much interest. There can be multiple examples of such anonymity requirements, including healthcare records, online shopping history, and movie ratings [20]. Electronic voting protocols are a prime example of a field where ensuring and verifying anonymity is crucial.

It is interesting to see how protocols are modelled symbolically for the analysis of such properties. In the Dolev-Yao model [10], one often requires special operators in order to capture certain behaviour. Many voting schemes employ an operation known as a *blind signature* [8]. A blind signature is one where the underlying object can be hidden (via a blinding factor), the now-hidden object signed, and then the blind removed to have the

signature percolate down to the underlying object. The FOO voting protocol given in [11] crucially uses blind signatures in order to obtain a signature on an encrypted object. [7] shows that the derivability problem for protocols involving blind signatures becomes DEXPTIME-hard. Protocols which do not use blind signatures often use homomorphic encryption or mix nets, which also make the modelling and verification quite complex [17].

Note that in most common models, terms are the only objects communicated. A "certificate" of an agent's validity – which is an intrinsically different object from a term containing an agent's vote, for example – is also modelled as a term in the term algebra. [4, 5], for example, augment the Dolev-Yao term syntax with an extra primitive ZK, which can be used to create a term that codes up a zero-knowledge proof. However, no direct logical inference is possible with these proof terms, and therefore, it is difficult to reason about what further knowledge agents can obtain using them. [22] proposes a departure from this paradigm, using *assertions* as a further abstraction that can be used for modelling protocols. Assertions, which code up certificates and have a separate proof system, can be sent by agents in addition to terms. The assertion algebra allows designers to model protocols involving certification in a more explicatory manner (by maintaining terms and certificates as separate objects). It also allows analysts to capture any increase in agents' knowledge achieved by deduction at the level of certificates.

So what are these assertions and how do they behave? Assertions include statements about various terms appearing in the protocol. These include instances of application-specific predicates and equalities between two different symbolic terms. Assertions can also be combined using the usual propositional connectives *and* ($\wedge$) and *or* ($\vee$). They also include a *says* operator, which works as an ownership mechanism for assertions, and disallows other agents from forwarding such an assertion in their own name. Perhaps the most crucial (and useful) addition to the assertion language here (over the system in [22]) is the existential quantifier. This allows us to quantify out any term from an assertion, thereby effectively hiding the actual term about which that assertion is made. Since existential assertions thus hide the private data used to generate a certificate, while revealing some partial information, they seem especially useful for capturing blinding (and similar operations with this goal) in voting protocols.

1.1 Related Work

Research on anonymity has been carried out for many years now. In the applied-pi calculus, [16] verifies anonymity for the FOO protocol, [2] studies general unlinkability and shows that this implies anonymity, and [19] provides an applied-pi based model incorporating aspects of the underlying communication mechanism (anonymous channels in particular).

There are also many epistemic logic-based approaches. [14] provides a logical framework built on modal epistemic logic for anonymity in multiagent systems; [12, 23] also define information-hiding properties in terms of agent knowledge; [15] provides a modular framework that allows one to analyze general unlinkability properties using function views, along with extensive case studies on anonymity and privacy.

Theorem provers have also been used to verify anonymity. [6] uses an automatic theorem

prover MCMAS for verification; [3] also specifies general unlinkability as an extension to the Inductive Method for security protocol verification in the theorem prover Isabelle.

In this paper, we extract a logical core of reasoning about certificates, translate the typical constructs used for voting protocols into proof properties, and employ equivalence-based reasoning for verifying anonymity. We also apply this technique to model two voting protocols, namely FOO and Helios, and to analyze anonymity for FOO.

2 Modelling the FOO protocol

2.1 Introduction to FOO

In [11], the authors introduce the FOO protocol for electronic voting, which has inspired many subsequent protocols. This protocol uses blinding functions and bit commitments in order to satisfy many desirable security properties, including anonymity. The voter V sends to the authority A his name, along with a blindsigned commitment to the vote v. The authority signs this term, and sends it back to V. V now unblinds this to obtain a signature on his commitment to the vote v, and sends that to the collector C. C adds the encrypted vote and V's commitment to the public bulletin board. V then sends to C the random bit r he used to create the vote commitment, so C can access the vote and update his tally. The protocol is presented in Figure 1a (see [11, 16] for a detailed explanation). Sends marked by $\looparrowright$ are over anonymous channels.

2.2 Modelling FOO with Assertions

In Figure 1b, we present the FOO protocol as modelled using assertions.

The voter V contacts the authority A with his vote v encrypted using a random key r_A. V also sends a certificate linking his name to his encrypted vote v. The V *says* prefix links V to the certificate about v, and thus informs the authority that V wishes to vote using the valid vote v, the encrypted form of which has been sent with the certificate. Note that this certificate automatically rules out replay attacks (of the kind where another agent V' copies V's published data off the bulletin board and replays it in her own name).

The authority A checks that the voter V has not voted earlier. If this check passes, A adds the fact that V has voted with the encrypted term $\{v\}_{r_A}$ to her database (so that V cannot vote again in the future) via an *insert* action. A then issues a certificate stating that V is a valid voter and wishes to vote with the encrypted term he sent A earlier, and that V claims that the term encrypted therein is a valid vote. The voter V now anonymously sends to the counter C the vote v encrypted in a new random key. This is accompanied by an existential assertion, which hides the voter's identity from C, while still convincing C that A has certified V and the sent vote to be valid.

We need three predicates here – `valid`, `elg`, and `voted`. The first two are predicates for stating the validity of the vote and the eligibility of the voter, respectively. The `voted` predicate is used for linking the voter and the vote. As we shall see in Section 4, we can add

$$
\begin{array}{lcl}
V \to A & : & V, [\text{blind}(\text{commit}(v, r), b)]_V \\
 & & \\
 & & \\
 & & \\
A \to V & : & [\text{blind}(\text{commit}(v, r), b)]_A \\
 & & \\
 & & \\
 & & \\
V \not\to C & : & [\text{commit}(v, r)]_A \\
C \to & : & \text{list}, [\text{commit}(v, r)]_A \\
V \to C & : & r
\end{array}
$$

(a) FOO Protocol with terms. $[x]_A$ denotes x signed by A.

$$
\begin{array}{lcl}
V \to A & : & \{v\}_{r_A}, V \ \textit{says} \ \{ \ \exists x, r : \{x\}_r = \{v\}_{r_A} \\
 & & \qquad\qquad \wedge \ \texttt{valid}(x) \ \ \} \\
A & : & \textit{deny} \ \ \exists x : \texttt{voted}(V, x) \\
A & : & \textit{insert} \ \texttt{voted}(V, \{v\}_{r_A}) \\
A \to V & : & A \ \textit{says} \\
 & & [\ \texttt{elg}(V) \wedge \texttt{voted}(V, \{v\}_{r_A}) \\
 & & \ \wedge \ V \ \textit{says} \ \{ \ \exists x, r : \{x\}_r = \{v\}_{r_A} \\
 & & \qquad\qquad \wedge \ \texttt{valid}(x) \ \} \ ] \\
V \not\to C & : & \{v\}_{r_C}, r_C, \\
 & & \exists X \ \exists y, s : A \ \textit{says} \\
 & & [\ \texttt{elg}(X) \wedge \ \texttt{voted}(X, \{y\}_s) \\
 & & \ \wedge \ X \ \textit{says} \ \{ \ \exists x, r : \{x\}_r = \{y\}_s \\
 & & \qquad\qquad \wedge \ \texttt{valid}(x) \ \} \ ] \\
 & & \wedge \ y = v
\end{array}
$$

(b) FOO Protocol with assertions.

Figure 1: FOO Protocol: Modelling with terms only and with assertions

such protocol-specific predicates to the assertion language in order to communicate succinct certificates (for example, here we use `valid`(v), instead of providing a disjunction over the finite set of valid votes for the value of v, which would grow longer as the set of allowable values grows larger).

3 Modelling Helios 2.0

3.1 Introduction to Helios

[1] introduces the voting scheme called Helios which has the desirable property of public auditability, i.e., even if Helios is fully corrupt, one can verify the integrity of an election outsourced to it. Helios provides unconditional integrity, while privacy is guaranteed if one trusts the Helios server, which doubles up as election administrator and trustee. The voter sends his vote to the Ballot Preparation System, which creates an encrypted ballot, which is then sealed and cast. The voter's identity and ballot are then posted on the public bulletin board. On closing the election, Helios removes voter names, shuffles all ballots, produces a proof of correct shuffling, and posts these on the board. After allowing some time for auditors to check the shuffling, Helios decrypts each ballot, produces a proof of correct decryption, and posts the tally on the bulletin board. Helios crucially uses auditing by various participants in order to guarantee correctness.

3.2 Helios 2.0

[9] demonstrates an attack on vote privacy in the basic Helios system in [1], where, by controlling more than half the voters, an adversary can get the compromised voters to copy a single (honest) voter's encrypted ballot off the bulletin board, and from the tally know whom that voter voted for. Note that this happens in spite of the Helios system itself being non-corrupt. In order to fix this, they introduce measures to weed out replayed ballots, and a linking mechanism between every ballot and the voter whose vote it is supposed to encrypt. They also replace the shuffling mechanism by a homomorphic encryption operation, and introduce trustees who are distinct from the election administrator. This introduces an extra assurance of vote privacy, since a corrupt administrator needs to corrupt some trustees in order to see a voter's unencrypted vote.

3.3 Modelling Helios 2.0 with Assertions

The voter first inputs his vote to a script which creates his ballot and sends it back to him with an assertion stating correctness. The voter can then choose to cast this vote, at which point the script submits his ballot and the assertion to the administrator. The administrator publishes the ballot and the assertion on the bulletin board. After some known deadline, the administrator homomorphically combines all ballots, and publishes the encrypted tally

$$
\begin{array}{rcl}
V \to S & : & v,\ V\ \mathit{says}\ \mathtt{valid}(v) \\
S \to V & : & b,\ S\ \mathit{says}\ \{\exists v : b = \mathit{ballot}(v) \wedge V\ \mathit{says}\ \mathtt{valid}(v)\} \\
V \to S & : & \mathit{cast} \\
S \to A & : & b,\ S\ \mathit{says}\ \{\exists v : b = \mathit{ballot}(v) \wedge V\ \mathit{says}\ \mathtt{valid}(v)\} \\
A & : & \mathit{deny}\ \mathtt{voted}(V) \\
A & : & \mathit{insert}\ \mathtt{voted}(V) \\
A \to BB & : & b,\ A\ \mathit{says}\ S\ \mathit{says}\ \{\exists v : b = \mathit{ballot}(v) \wedge V\ \mathit{says}\ \mathtt{valid}(v)\} \\
\text{Suppose} & b_1, \ldots, b_k & \text{were the ballots cast and published on the bulletin board.} \\
A \to BB & : & t,\ A\ \mathit{says}\ [\exists s : t = \mathit{ballot}(s) \wedge \\
 & & \{\exists v_1, \ldots, v_k : s = \mathit{sum}(v_1, \ldots, v_k) \wedge \bigwedge_{i=1}^{k} b_i = \mathit{ballot}(v_i)\}]
\end{array}
$$

Figure 2: Helios 2.0 Protocol with assertions

along with an assertion stating correctness of the tally. The trustees can then decrypt this tally, and the administrator publishes the result.

In Figure 2 we model Helios 2.0 with assertions. We do not include the final step, where the trustees decrypt the final encrypted tally and publish it onto the bulletin board. Note that this model, much like the terms-only model in [9], requires us to add a homomorphic encryption operation to our term algebra. However, we can incorporate the weeding out of replayed ballots and establishing the link between ballots and voters by the use of assertions alone, instead of having to send extra terms. Note that in order for an agent V_2 to copy V_1's vote and replay it to A, V_2 would need to make an assertion of the form $S\ \mathit{says}\ \{\exists v : b = \mathit{ballot}(v) \wedge V_2\ \mathit{says}\ \mathtt{valid}(v)\}$, which would contradict the sending in V_1's name. Thus we can establish a link between vote and voter, while also disallowing replays. We merely need to add a homomorphic encryption operation to the term algebra, since our assertions, as of now, are not capable of capturing this operation.

4 Assertions: Theory

We fix the following countable sets – a set $\mathscr{V}$ of variables, a set Ag of agents, a set $\mathscr{N}$ of nonces, and a set of $\mathscr{K}$ of keys. We assume that every $k \in \mathscr{K}$ has an inverse key, denoted $\mathit{inv}(k)$. The set of basic terms $\mathscr{B}$ is defined to be $Ag \cup \mathscr{N} \cup \mathscr{K}$. The set of terms $\mathscr{T}$ is given by the following syntax:

$$t := m \mid (t_1, t_2) \mid \{t\}_k$$

where $m \in \mathscr{B} \cup \mathscr{V}$, and $k \in \mathscr{K} \cup \mathscr{V}$. A term with no variables occurring in it is called a *ground term.*

$\dfrac{}{X \cup \{t\} \vdash t}\ ax$	
$\dfrac{X \vdash t_1 \quad X \vdash t_2}{X \vdash (t_1, t_2)}\ pair$	$\dfrac{X \vdash (t_1, t_2)}{X \vdash t_i}\ split$
$\dfrac{X \vdash t \quad X \vdash k}{X \vdash \{t\}_k}\ enc$	$\dfrac{X \vdash \{t\}_k \quad X \vdash inv(k)}{X \vdash t}\ dec$

Table 1: The Dolev-Yao derivation system

The system of rules for deriving new ground terms from old is given in Table 1. The rules are presented in terms of sequents $X \vdash t$ where X is a finite set of ground terms, and t is a ground term.

4.1 Assertions and derivations

We now present the formal details of the model with assertions, a version of which was first proposed in [22]. The set of assertions, $\mathscr{A}$, is given by the following syntax (fixing a set of variables, and a set of predicates for each arity):

$$\begin{aligned}\alpha := t = t' \ &\mid\ \alpha_1 \vee \alpha_2 \ \mid\ \alpha_1 \wedge \alpha_2 \ \mid\ \exists x : \alpha \ \mid\ m \ says\ \alpha \\ &\mid\ \texttt{valid}(m) \ \mid\ \texttt{elg}(m) \ \mid\ \ldots \ \mid\ m \ sent\ t \ \mid\ m \ sent\ \alpha\end{aligned}$$

where $t \in \mathscr{T}$, $m \in Ag \cup \mathscr{V}$, and `valid` and `elg` are application-specific predicates. The ellipses signify that one may add more such simple predicates, depending on the application requirements (as in the FOO protocol, from Section 2.2). A *ground assertion* is one with no free variables.

The set of assertions is a positive fragment of existential first-order logic. The intention is that in addition to ground terms, agents also communicate ground assertions to each other. Agents are allowed to assert equality of terms, and basic predicates on terms, as well as disjunctions and conjunctions. They can also "sign" assertions by use of the *says* operator. They also have the capability of existentially abstracting some terms from an assertion, thereby modelling *witness hiding*. The sole use of the *sent* operator is to enable an observer to record who communicated a term or an assertion.

In the course of participating in a protocol, agents accumulate a database of ground terms and ground assertions communicated to them. The proof system for assertions is presented in Table 2. The rules are presented in terms of sequents $X, \Phi \vdash \alpha$, where X is a finite set of ground terms and Φ is a finite set of assertions (which are not necessarily ground).

Equality assertions form a central part of communications between agents. Note that an agent A can derive $t = t$ only when all basic subterms of t can be derived by A. The recipient of an equality assertion can use the rules provided in Table 2 to reason further about the terms involved therein. Our rules for equality are fairly intuitive and reflect basic properties

$\dfrac{X \vdash_{dy} m}{X, \Phi \vdash m = m}\ [m \in \mathscr{B} \cup \mathscr{V}]$	$\dfrac{}{X, \Phi \cup \{\alpha\} \vdash \alpha}\ ax$	$\dfrac{X, \Phi \vdash \alpha(t)\ \ X, \Phi \vdash t = t'}{X, \Phi \vdash \alpha(t')}$
$\dfrac{X, \Phi \vdash s = t\ \ X, \Phi \vdash t = u}{X, \Phi \vdash s = u}$	$\dfrac{X, \Phi \vdash s = t}{X, \Phi \vdash t = s}$	$\dfrac{X, \Phi \vdash (s_0, s_1) = (t_0, t_1)}{X, \Phi \vdash s_i = t_i}$
$\dfrac{X, \Phi \vdash s = s'\ \ X, \Phi \vdash t = t'}{X, \Phi \vdash (s, t) = (s', t')}$	$\dfrac{X, \Phi \vdash \{s_0\}_{s_1} = \{t_0\}_{t_1}}{X, \Phi \vdash s_i = t_i}\ \bullet$	$\dfrac{X, \Phi \vdash s = s'\ \ X, \Phi \vdash m = m'}{X, \Phi \vdash \{s\}_m = \{s'\}_{m'}}$
$\dfrac{X, \Phi \vdash m = n}{X, \Phi \vdash \alpha}\ \bot\ [m, n \in \mathscr{B}, m \neq n]$		$\dfrac{X, \Phi \vdash \alpha\ \ X \vdash_{dy} sk(A)}{X, \Phi \vdash A\ says\ \alpha}\ says_A$
$\dfrac{X, \Phi \vdash \alpha\ \ X, \Phi \vdash \beta}{X, \Phi \vdash \alpha \wedge \beta}\ \wedge i$	$\dfrac{X, \Phi \vdash \alpha_1 \wedge \alpha_2}{X, \Phi \vdash \alpha_i}\ \wedge e$	$\dfrac{X, \Phi \vdash A\ says\ \alpha}{X, \Phi \vdash \alpha}\ strip$
$\dfrac{X, \Phi \vdash \alpha_i}{X, \Phi \vdash \alpha_1 \vee \alpha_2}\ \vee i$	$\dfrac{X, \Phi \vdash \alpha \vee \beta\ \ X, \Phi \cup \{\alpha\} \vdash \delta\ \ X, \Phi \cup \{\beta\} \vdash \delta}{X, \Phi \vdash \delta}\ \vee e$	
$\dfrac{X, \Phi \vdash \alpha(t)}{X, \Phi \vdash \exists x : \alpha(x)}\ \exists i$	$\dfrac{X, \Phi \vdash \exists x : \alpha(x)\ \ X, \Phi \cup \{\alpha(y)\} \vdash \beta}{X, \Phi \vdash \beta}\ \exists e$	

Table 2: Derivation rules for assertions. We assume that $X \vdash_{dy} x$ for all variables x, and that $inv(x) = x$. In the $\bullet$ rule, we require that $X \vdash_{dy} inv(s_1)$ and $X \vdash_{dy} inv(t_1)$. In the $\exists e$ rule, we require that $y \notin Vars(X, \Phi \cup \{\beta\})$.

of the pairing and encryption operations. Equality assertions are most likely to be used in existentially quantified assertions. Notable among the other rules are $says_A$, which allows the possessor of $sk(A)$ to "sign" an assertion in A's name, and *strip*, which allows one to strip the sign in $A\ says\ \alpha$ and use α in local reasoning.

These rules allow agents to carry out non-trivial inferences, potentially learning more than was intended by the protocol. Suppose an agent A has a term $\{v\}_k$, which he knows be a nonce encrypted with some key, but whose inverse he does not have access to. One would presume that A therefore should have no idea about the value of v. However, it is possible for assertions about $\{v\}_k$ to reveal more information to A. Suppose A manages to obtain two certificates $\exists x, y : \{v\}_k = \{x\}_y\ \wedge\ (x = 0 \vee x = 1)$ and $\exists x, y : \{v\}_k = \{x\}_y \wedge (x = 0 \vee x = 2)$. Let us call these assertions $\exists x, y : \alpha(x, y)$ and $\exists x, y : \alpha'(x, y)$. These two assertions are in A's database of assertions Φ. Let a, b, a', b' be new variables that do not occur in Φ. Consider $\Phi \cup \{\alpha(a, b), \alpha'(a', b')\}$. From $\{v\}_k = \{a\}_b$ and $\{v\}_k = \{a'\}_{b'}$, we get $\{a\}_b = \{a'\}_{b'}$, and hence $a = a'$ and $b = b'$. From the other parts of α and α', and using transitivity, we get $a = 0 \vee a = 1$ and $a = 0 \vee a = 2$. We use disjunction elimination to get $a = 0$. From this we conclude that $\{v\}_k = \{0\}_b$, and hence $\Phi \cup \{\alpha(a, b), \alpha'(a', b')\} \vdash \exists y : (\{v\}_k = \{0\}_y)$.

Therefore, using the $\exists e$ rule, we get $\Phi \vdash \exists y : (\{v\}_k = \{0\}_y)$.

In the formal model of [4, 5], each zkp term proves a formula involving some private and some public variables. The recipient of a zkp term is deemed to have knowledge of the terms used in place of the public variables, but not the private ones. We adopt a similar convention. For an assertion α, if an equality of the form $t = t'$ occurs in it, or if α involves the application of a predicate to a term t, then α reveals t. However, if a term of the form $\{v\}_k$, say, appears in α, then α does not reveal v. We also adopt the convention that every term revealed by an assertion is sent earlier in the protocol.

4.2 Actions, roles and protocols

There are six type of actions – send, anonymous send, receive, confirm, deny, and insert. Sends, anonymous sends, and receives are of the form $+A$: $(\vec{m})(t, \alpha)$, $+A^*$: $(\vec{m})(t, \alpha)$ and $-A$: (t, α) respectively, where $A \in Ag \cup \{id\}$ (where id is a dedicated variable that stands for the agent performing the action), $\vec{m} \subseteq \mathscr{V} \cup \mathscr{N} \cup \mathscr{K}$ stands for nonces and keys that are *fresh* which should be instantiated with hitherto unused values in each occurrence of this action, $t \in \mathscr{T}$ and $\alpha \in \mathscr{A}$. The A: *confirm* α and A: *deny* α actions allow A to branch on whether or not he can derive α, while A: *insert* α allows A to add previously unknown true assertions into her database. For $A \in Ag \cup \{id\}$, an A-action is an action which involves A. A *ground action* is one without any variable occurrence. An A-role is a finite sequence of A-actions. A role is an A-role for some $A \in Ag \cup \{id\}$. A *protocol* Pr is a finite set of roles.

Given a sequence of actions $\eta = a_1 \cdots a_n$, we say that the variable x *originates at* i if x occurs in a_i and does not occur in a_j for any $j < i$. A variable x occurring in a role η is said to be *bound* if it originates at i and either a_i is a receive action, or $a_i = +A$: $(\vec{y})(t, \alpha)$ is a send action with $x \in \vec{y}$.

As an example, we show the voter role for the FOO protocol from Section 2. In this role, v and id stand for the vote and voter respectively, while k, k' are fresh keys, and *auth* is a bound variable (since it originates in a receive) which stands for the authority with whom the voter interacts.

$$
\begin{array}{rcl}
+id & : & (k)\ \{v\}_k, id\ says\ \{\exists x, r : \{x\}_r = \{v\}_k \wedge \mathtt{valid}(x)\} \\
-id & : & auth\ says\ [\mathtt{elg}(id)\ \wedge\ \mathtt{voted}(id, \{v\}_k) \\
 & & \quad \wedge\ id\ says\ \{\ \exists x, r : \{x\}_r = \{v\}_k \wedge\ \mathtt{valid}(x)\ \}\ ] \\
+id^* & : & (k')\ (\{v\}_{k'}, k'), \\
 & & \exists X, y, s : auth\ says\ [\mathtt{elg}(X) \wedge \mathtt{voted}(X, \{y\}_s) \\
 & & \wedge\ X\ says\ \{\exists x, r : \{x\}_r = \{y\}_s \wedge \mathtt{valid}(x)\}] \wedge\ y = v
\end{array}
$$

The authority and counter roles can also be extracted from the protocol description in a similar manner.

4.3 Runs of a protocol

Even though the roles of a protocol mention variables, its *runs* (or executions) consist only of ground terms and assertions exchanged in various *instances* of the roles. An instance of a role is formally specified by a *substitution* σ, which is a partial map from $\mathcal{V}$ to the set of all ground terms. We lift σ for terms, assertions and actions in the standard manner. σ is said to be *suitable* for an action a if $\sigma(a)$ is an action, i.e. a typing discipline is followed. A substitution is suitable for a role η if it is defined on all free variables of η and suitable for all actions in η.

A *session* of a protocol Pr is a sequence of actions of the form $\sigma(\eta)$, where $\eta \in Pr$ and σ is suitable for η.

A run of a protocol is an interleaving of sessions in which each agent can construct the messages that it communicates. This is formalized by a notion of *knowledge state*, which represents all the terms and assertions that each agent knows. A *control state* is a record of progress made by an agent in the various sessions he/she participates in.

A knowledge state ks is a tuple $((X_A, \Phi_A)_{A \in Ag})$, where X_A (resp. Φ_A) is the set of ground terms (resp. ground assertions) belonging to an agent A. A control state S is a finite set of sequences of actions. A protocol state is a pair (ks, S) where ks is a knowledge state and S is a control state.

Definition 1. *Let (ks, S) and (ks', S') be two states of a protocol Pr, and let b be a ground action. We say that $(ks, S) \xrightarrow{b} (ks', S')$ iff there is a session $\eta = a \cdot \eta' \in S$ and a substitution σ suitable for η' such that:*

- $b = \sigma(a)$
- $S' = (S \setminus \{\eta\}) \cup \{\sigma(\eta')\}$
- $ks \xrightarrow{b} ks'$ *as given in Table 3.*

In Definition 1, we add $\sigma(\eta')$ rather than η', in order to update the substitution associated with the session on executing the action. This update reflects the new values generated for each fresh nonce variable (in case the action is a send) or the new bindings for input variables (in case the action is a receive). For instance, if $\eta = a \cdot \eta'$ where $a = -A\colon ((x, y), \alpha(x, y))$ and $b = -A\colon ((t, t'), \alpha(t, t'))$, then $\sigma = [x := t, y := t']$. Any occurrence of x in η' is bound to t.

Note the crucial difference between the updates for sends and anonymous sends – in the former, the intruder updates its state with A *sent* t and A *sent* α, whereas in the latter, no sender information is available to any observer (including the intruder).

An *initial control state* of Pr is a finite set of sessions of Pr. In the *initial knowledge state*, each agent has her own secret keys and shared keys, all public keys in her database, and potentially some constants of Pr.

Definition 2. *A run of a protocol Pr is $(ks_0, a_1 \cdots a_n)$ such that ks_0 is an initial knowledge state, and there exist sequences $ks_1, \ldots, ks_n$ and $S_0, \ldots, S_n$ such that $(ks_{i-1}, S_{i-1}) \xrightarrow{a_i} (ks_i, S_i)$ for all $i \leq n$.*

Action b	Enabling conditions	Updates
$+A$: $(\vec{m})(t,\alpha)$	$X_A \cup \vec{m} \vdash t$ $X_A \cup \vec{m}, \Phi_A \vdash \alpha$	$X'_A = X_A \cup \vec{m} \quad X'_I = X_I \cup \{t\}$ $\Phi'_I = \Phi_I \cup \{\alpha, A \textit{ sent } t, A \textit{ sent } \alpha\}$
$+A^*$: $(\vec{m})(t,\alpha)$	$X_A \cup \vec{m} \vdash t$ $X_A \cup \vec{m}, \Phi_A \vdash \alpha$	$X'_A = X_A \cup \vec{m} \quad X'_I = X_I \cup \{t\}$ $\Phi'_I = \Phi_I \cup \{\alpha\}$
$-A$: (t,α)	$X_I \vdash t$ $X_I, \Phi_I \vdash \alpha$	$X'_A = X_A \cup \{t\}$ $\Phi'_A = \Phi_A \cup \{\alpha\}$
A: *confirm* α	$X_A, \Phi_A \vdash \alpha$	No change
A: *deny* α	$X_A, \Phi_A \nvdash \alpha$	No change
A: *insert* α	Always enabled	$\Phi'_A = \Phi_A \cup \{\alpha\}$

Table 3: Enabling conditions for $ks \xrightarrow{b} ks'$. We assume that for each agent A, (X_A, Φ_A) and (X'_A, Φ'_A) represent A's knowledge in ks and ks', respectively.

4.4 Notes on implementability

A central aspect of this model is that communicated assertions are "believed" by the recipients. This is reflected in the updates for receive actions. On the other hand, it is not possible for a malicious agent to inject "falsehoods" into the system, as evidenced by the enabling conditions which only allow derivable assertions to be communicated. How might all this be realized in practice?

An implementation is to demand that every communicated assertion be translated into an appropriate zero knowledge proof. But suppose an agent receives ZKPs for assertions α and β from A and B, and wishes to send $\alpha \wedge \beta$ to someone else. For this, she should have the capacity to produce a ZKP for $\alpha \wedge \beta$. This implements the $\wedge i$ rule in our system. Clearly this requires some mechanism for composing ZKPs. Such a system has been studied in [18], which proposes a logical language close to ours, and also discusses modular construction of ZKPs, based on the seminal work on composability of ZKPs [13].

However, [18] has some restrictions on the proof rules for which one can modularly construct ZKPs. For instance, they do not consider disjunction elimination or existential elimination. Nevertheless, we consider these rules since they are at the heart of potential attacks (as illustrated by the earlier example). This situation can be handled formally by making a distinction between rules that are "safe for composition" and rules that are not. A rule like $\wedge i$ is safe for composition, for example, whereas $\vee e$ might not be. We then adopt the restriction that we communicate assertions that are derived using only safe rules. If the derivation of an assertion necessarily involves unsafe rules, then it cannot be communicated to another agent, even though this derivation itself is allowed for local reasoning. In this paper, we therefore consider both local reasoning to derive more assertions (to gain more knowledge about some secrets, for instance) as well as deriving communicable assertions.

5 Formalizing anonymity

Informally, we say that a voting protocol satisfies anonymity if in all executions of the protocol, no adversary can deduce the connection between a voter and her vote. One way to formalise it is to consider a run ρ where voter V_0 voted 0 and voter V_1 voted 1, and show that there is some run ρ' where the votes of V_0 and V_1 are swapped and every other voter acts the same as in ρ, such that even the most powerful intruder I (who has access to all keys of the authorities) cannot distinguish ρ from ρ'.

Definition 3. *Let (ks, ρ) and (ks', ρ') be two runs of Pr, where $\rho = a_1 \cdots a_n$ and $\rho' = a'_1 \cdots a'_n$. Let t_i and t'_i be the terms communicated in a_i and a'_i, respectively. Let (X, Φ) and (X', Φ') be the knowledge states of I at the end of each run.*

We say that (ks, ρ) is I-indistinguishable from (ks', ρ') – denoted $(ks, \rho) \sim_I (ks', \rho')$ – if for all assertions $\alpha(x_1, \ldots, x_k)$ and all sequences $i_1 < \cdots < i_k \leq n$:

$$X, \Phi \vdash \alpha(t_{i_1}, \ldots, t_{i_k}) \text{ iff } X', \Phi' \vdash \alpha(t'_{i_1}, \ldots, t'_{i_k}).$$

One can view the parameters $x_1, \ldots, x_k$ occurring in the above definition as *handles*, and the mapping from $x_1, \ldots, x_k$ to $t_{i_1}, \ldots, t_{i_k}$ as an *active substitution*. Parametrized assertions $\alpha(x_1, \ldots, x_k)$ constitute *tests* on each run of the protocol. Thus the above notion is related to the notion of static equivalence that is central to protocol modelling in the applied-pi calculus [4, 5, 16]. Note that the notion of indistinguishability we use here is trace-based, as that fits naturally with our model. But it is also possible to have a bisimulation-based definition, and adapt our proof ideas.

Consider a voting protocol Pr with three roles – *voter*, *authority* and *counter*, and two phases: *authorization* and *voting*. For simplicity, we assume that there are two fixed agents A and C who play the authority and counter role, respectively. If there is only one voter in a run, then obviously his/her vote can be linked to him/her. If a voter's vote is counted during the authorization phase, then we might have a situation where a vote is cast by a voter before anyone else has been authorized. This again is an easy violation of anonymity. Therefore we assume that in any run of Pr, there are at least two agents playing the voter role, and all $V_i \to A$ actions precede all $V_j \to C$ actions.

Fix voter names V_0, V_1, and votes v_0 and v_1. A session η of Pr is said to be an (i, j)-session if η maps id to V_i and v to v_j.

Definition 4. *We say that Pr satisfies anonymity if for every initial knowledge state $ks = (X, \Phi)$ such that $X_A \cup X_C \subseteq X_I$, and for every run (ks, ρ) which includes a $(0, 0)$-session and a $(1, 1)$-session, there is a run (ks, ρ') which includes a $(1, 0)$-session and a $(0, 1)$-session such that $(ks, \rho) \sim_I (ks, \rho')$.*

Theorem 5. *The FOO protocol satisfies anonymity.*

Proof. Recall the voter role for FOO from Section 4.2. Consider a run (ks, ρ) of FOO whose initial control state is $S \cup \{\eta_0, \eta_1\}$, where η_0 is the $(0, 0)$-session and η_1 is the $(1, 1)$-session. Let η_2 and η_3 be the $(0, 1)$-session and $(1, 0)$-session, respectively. We construct a run ρ'

which includes η_2 and η_3 such that $(ks, \rho) \sim_I (ks, \rho')$. The session η_0 assigns values p and r to the keys k and k' from the role description, while η_1 assigns values q and s respectively. For ease of notation, we denote v_0 and v_1 by u and v respectively, and $d = \{u\}_p$ and $e = \{v\}_q$.

Suppose $\rho = a_1 \cdots a_n$. Assume without loss of generality that both sessions η_0 and η_1 are fully played out in ρ. Also without loss of generality, let $i < j < k < l$ be indices such that the send actions of η_0 are a_i and a_k, and the send actions of η_1 are a_j and a_l, where

$$a_i = +V_0\colon (p)(d, \beta(d)) \quad \text{and} \quad a_k = +V_0^*\colon (r)((\{u\}_r, r), \gamma(u))$$
$$a_j = +V_1\colon (q)(e, \beta(e)) \quad \text{and} \quad a_l = +V_1^*\colon (s)((\{v\}_s, s), \gamma(v))$$

We build $\rho' = b_1 \cdots b_n$ as shown in Figure 3.

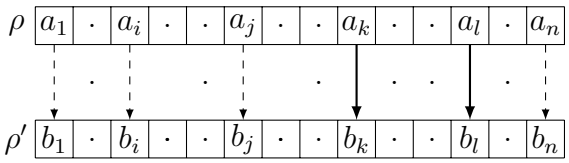

Figure 3: Building ρ' from ρ. The dashed arrows capture $b_m = a_m[d \mapsto e, e \mapsto d]$, for all $m \notin \{l, k\}$. For $m \in \{l, k\}$, the thick arrows stand for $b_m = a_m[V_0 \mapsto V_1, V_1 \mapsto V_0]$.

Observe that ρ' is also a run of FOO starting from the state $(ks, S \cup \{\eta_2, \eta_3\})$, where η_2 contains b_i and b_l, and η_3 contains b_j and b_k. We crucially use the fact that we do not fix the instances of the fresh nonces a priori, so we can swap the action containing p as a fresh nonce with the one containing q as a fresh nonce, for example.

For any term t (resp. assertion α), we define $\mathsf{swp}(t)$ (resp. $\mathsf{swp}(\alpha)$) to be the result of changing all occurrences of d to e and vice versa. swp is lifted to sets of terms and assertions as usual.

Let (X, Φ) and (X', Φ') be the knowledge states of I at the end of ρ and ρ' respectively. It is evident from the construction of ρ' that $X' = \mathsf{swp}(X)$. Furthermore, it is easy to see that neither X nor X' derive either p or q, and that $X \vdash_{dy} t$ iff $X' \vdash_{dy} \mathsf{swp}(t)$.

It can also be seen that $\Phi' = \mathsf{swp}(\Phi)$, as elaborated below. For every m, if a_m communicates α, then b_m communicates $\mathsf{swp}(\alpha)$. The other formulas added to Φ are *sent* assertions. For every action a_m other than a_k and a_l, the sender of b_m is unchanged from a_m. Therefore, a *sent* assertion with the same sender name would be added to Φ and Φ'. For a_k and a_l, no *sent* assertions are added since these are anonymous sends. Therefore, $\Phi' = \mathsf{swp}(\Phi)$.

We now prove that $X, \Phi \vdash \alpha(t_{i_1}, \ldots, t_{i_k})$ iff $X', \Phi' \vdash \alpha(t'_{i_1}, \ldots, t'_{i_k})$, for all assertions $\alpha(x_1, \ldots, x_k)$. It suffices to prove that $X, \Phi \vdash \alpha$ iff $X', \Phi' \vdash \mathsf{swp}(\alpha)$ for all α. For every $\exists : \delta$, let y_δ be a variable that does not occur in Φ. A set Θ is said to be *closed under witnesses* if $\delta(y_\delta) \in \Theta$ for all $\exists y : \delta \in \Theta$. Let Π be the smallest superset of Φ closed under witnesses. We use Π' to denote $\mathsf{swp}(\Pi)$. It can be shown by an analysis of derivations that $X, \Phi \vdash \alpha$ iff $X, \Pi \vdash_1 \alpha$ and $X', \Phi' \vdash \alpha$ iff $X', \Pi' \vdash_1 \alpha$, where $\vdash_1$ denotes derivability without using the $\exists e$ rule. Note that both X, Π and X', Π' are safe for d and e in the following sense. They do not derive equalities of the form $p = t$ or $q = t$ for any term t, and they do not derive

equalities of the form $d = t'$ or $e = t'$ where t' is a term containing a non-variable. We now prove the final claim needed for indistinguishability of ρ and ρ'.

Claim. For any α, $X, \Pi \vdash_1 \alpha$ iff $X', \Pi' \vdash_1 \mathsf{swp}(\alpha)$.

Proof of Claim We prove the implication from left to right, by induction on structure of derivations. The other direction holds by symmetry. Suppose π is a derivation of $X, \Pi \vdash \alpha$, with last rule r.

$r = \boldsymbol{ax}$**:** Suppose $\alpha \in \Pi$. It follows that $\mathsf{swp}(\alpha) \in \Pi'$.

r **is equality of encrypted terms:** π looks as follows.

$$\frac{\begin{array}{cc}\overset{\pi_0}{\vdots} & \overset{\pi_1}{\vdots} \\ X, \Pi \vdash s = s' & X, \Pi \vdash m = m'\end{array}}{X, \Pi \vdash \{s\}_m = \{s'\}_{m'}}$$

Suppose $\{s\}_m$ is either d or e. Then m is either p or q, and this would mean that $p = m'$ or $q = m'$ is derivable, contradicting safety of X, Π. Therefore $\{s\}_m$ is not equal to either d or e. By induction hypothesis, $\mathsf{swp}(s = s')$ is derivable from X', Π', and hence $\mathsf{swp}(\{s\}_m = \{s'\}_{m'})$ is also derivable.

r **is equality of decrypted terms:** In this case, π is of the following form

$$\frac{\begin{array}{ccc}\overset{\pi_0}{\vdots} & \overset{\pi_1}{\vdots} & \overset{\pi_2}{\vdots} \\ X, \Pi \vdash \{s\}_m = \{s'\}_{m'} & X \vdash_{dy} inv(m) & X \vdash_{dy} inv(m')\end{array}}{X, \Pi \vdash s = s'}$$

By induction hypothesis, it follows that $X', \Pi' \vdash \mathsf{swp}(\{s\}_m) = \mathsf{swp}(\{s'\}_{m'})$. Observe that neither $\{s\}_m$ nor $\{s'\}_{m'}$ is the same as d or e (for otherwise we would have that $X \vdash_{dy} p$ or $X \vdash_{dy} q$, which is an impossibility). Thus any occurrence of d or e in $\{s\}_m$ is inside s, and similarly for $\{s'\}_{m'}$. Thus $\mathsf{swp}(\{s\}_m) = \{\mathsf{swp}(s)\}_m$ and $\mathsf{swp}(\{s'\}_{m'}) = \{\mathsf{swp}(s')\}_{m'}$. Therefore $\mathsf{swp}(s) = \mathsf{swp}(s')$ is also derivable. ($inv(m)$ and $inv(m')$ are derivable from X' since they are derivable from X and do not mention d or e.)

The rest of the cases are on similar lines (or simpler, appealing to the induction hypothesis). □

6 Conclusion

In this paper, we extended the model of [22] by adding exisential assertions to the language, as a tool to hide private data used to generate certificates. These assertions are especially useful in coding up constructs that are common to voting protocols. We showed how to

specify protocols in this model, and formalised the notion of anonymity in terms of indistinguishability. In a non-trivial example of analysis in our model, we proved anonymity for the FOO protocol.

One way of extending this model is by adding a background theory of universally quantified sentences. Such a theory is a standard part of many authorization systems. For instance, if an agent A communicates to B the assertion $\exists x : \texttt{voted}(V, x)$ and if the background theory contains the assertion

$$\forall X, x : \{\texttt{voted}(X, x) \Rightarrow \texttt{elg}(X)\}$$

then B can conclude $\texttt{elg}(V)$. More detailed examples are found in [4, 18]. It is an important ingredient in many systems, and we can easily incorporate it in our theoretical model.

References

[1] B. Adida. Helios: Web-Based Open-Audit Voting. In Proc. of the 17th conference on Security symposium (SS'08), pages 335–348, 2008.

[2] M. Arapinis, T. Chothia, E. Ritter, and M. Ryan. Analysing Unlinkability and Anonymity using the Applied Pi Calculus. In 23rd IEEE Computer Security Foundations Symposium, pages 107–121, 2010.

[3] D. Butin, D. Gray, and G. Bella. Towards Verifying Voter Privacy Through Unlinkability. In *Proc. ESSoS13, LNCS*, pages 91–106, 2013.

[4] M. Backes, C. Hritcu, and M. Maffei. Type-Checking Zero-Knowledge. In *Proc. 15th ACM CCS*, pages 357–370, 2008.

[5] M. Backes, M. Maffei and D. Unruh. Zero-Knowledge in the Applied Pi-calculus and Automated Verification of the Direct Anonymous Attestation Protocol. In *IEEE Symposium on Security and Privacy*, pages 202–215, 2008.

[6] I. Boureanu, A. V. Jones, and A. Lomuscio. Automatic Verification of Epistemic Specifications under Convergent Equational Theories. In *Proc. 11th AAMAS*, pages 1141–1148, 2012.

[7] A. Baskar, R. Ramanujam, and S. P. Suresh. A dexptime-Complete Dolev-Yao Theory with Distributive Encryption. In *Proc. Mathematical Foundations of Computer Science LNCS 6281*, pages 102–113, 2010.

[8] D. Chaum. Blind Signatures for Untraceable Payments. In *Advances in Cryptology*, pages 199–203, 1983.

[9] V. Cortier, and B. Smyth. Attacking and Fixing Helios: An Analysis of Ballot Secrecy. In *Proc. Computer Security Foundations Symposium*, pages 297–311, 2011.

[10] D. Dolev, and A. Yao. On the security of public key protocols. In *IEEE Transactions on Information Theory*, 198–208, 1983.

[11] A. Fujioka, T. Okamoto, and K. Ohta. A Practical Secret Voting Scheme for Large Scale Elections. In *International Workshop on the Theory and Application of Cryptographic Techniques*, pages 244–251, 1992.

[12] J.W. Gray, and P.F. Syverson. A Logical Approach to Multilevel Security of Probabilistic Systems. In *Distributed Computing, 11*, pages 73–90, 1998.

[13] J. Groth, and A. Sahai. Efficient Non-interactive Proof Systems for Bilinear Groups. In *Advances in Cryptology – EUROCRYPT 2008*, pages 415–432, 2008.

[14] J. Y. Halpern, and K. R. O'Neill. Anonymity and Information Hiding in Multiagent Systems. In *Journal of Computer Security, 13(3)*, pages 483–514, 2005.

[15] D. Hughes, and V. Shmatikov. Information hiding, Anonymity and Privacy: A Modular Approach. In *Journal of Computer security 12(1)*, pages 3-36, 2004.

[16] S. Kremer, and M. Ryan. Analysis of an Electronic Voting Protocol in the Applied Pi Calculus. In *European Symposium on Programming*, pages 186–200, 2005.

[17] Pascal Lafourcade, Denis Lugiez, and Ralf Treinen. Intruder Deduction for the Equational Theory of Abelian Groups with Distributive Encryption. *Information and Computation*, 205(4):581–623, April 2007.

[18] M. Maffei, K. Pecina, and M. Reinert. Security and Privacy by Declarative Design. In *IEEE 26th CSF Symposium*, pages 81–96, 2013.

[19] S. Mauw, J. Verschuren, and E. P. de Vink. Data Anonymity in the FOO Voting Scheme. In *Electronic Notes in Theoretical Computer Science (168)*, pages 5–28, 2007.

[20] The Netflix Prize. `http://www.netflixprize.com/index.html`

[21] A. J. Paverd, A. Martin, and I. Brown. Modelling and Automatically Analysing Privacy Properties for Honest-but-Curious Adversaries. *Technical Report*, `https://www.cs.ox.ac.uk/people/andrew.paverd/casper/casper-privacy-report.pdf`, 2014.

[22] R. Ramanujam, V. Sundararajan, and S. P. Suresh. Extending Dolev-Yao with Assertions. In *Proc. ICISS, LNCS 8880*, pages 50–68, 2014.

[23] P.F. Syverson, and S.G. Stubblebine. Group Principals and the Formalization of Anonymity. In *Proc. FM, LNCS 1708*, pages 814–833, 1999.